How Bandwidth Affects the *CONGEST* Model

Dennis Olivetti[*1,2]

1 **Gran Sasso Science Institute, L'Aquila, Italy.**
2 **Institut de Recherche en Informatique Fondamentale CNRS and University Paris Diderot, France.**

arXiv:1704.06092v1 [cs.DC] 20 Apr 2017

Abstract

The CONGEST model for distributed network computing is well suited for analyzing the impact of limiting the throughput of a network on its capacity to solve tasks efficiently. For many "global" problems there exists a lower bound of $\Omega(D+\sqrt{n/B})$, where B is the amount of bits that can be exchanged between two nodes in one round of communication, n is the number of nodes and D is the diameter of the graph. Typically, upper bounds are given only for the case $B = O(\log n)$, or for the case $B = +\infty$.

In this work we study tradeoffs between the round complexity for solving a task, and the bandwidth of the links. We show three different classes of problems, namely, *bandwidth efficient*, *bandwidth sensitive* and *bandwidth insensitive*. The first is the class of problems having a round complexity that fully scales with the bandwidth. We show that the All Pairs Shortest Path problem belongs to this class. The second class includes all problems having a round complexity that scales with the bandwidth, but not linearly. This is the case of the Minimum Spanning Tree construction and the Single Source Shortest Path problems. At last, the third class contains problems whose complexity does not depend on the size of the messages. We show that this class is not empty.

1 Introduction

1.1 Context

The links of networks typically have limited bandwidth. The CONGEST model for distributed network computing captures this constraint. In this model, an algorithm proceeds in synchronous rounds. At each round, every node can send B bits to each of its neighbors in the network (these B bits do not need to be the same for all neighbors). A typical value for B is $O(\log n)$ in n-node networks. This value is sufficient to transmit an integer in a polynomial range, like the identifier of a node, or the weight of an edge.

One celebrated result in this context is a minimum-weight spanning tree (MST) construction algorithm that performs in $O(D + \sqrt{n}\log^* n)$ rounds in diameter-D n-node networks [12]. This complexity is optimal for $B = O(\log n)$ [18]. On the other hand, all (computable) tasks can be solved in $O(D)$ rounds in diameter-D networks whenever there is no limitation on the bandwidth, by gathering all data at one node, computing the solution at that node, and broadcasting that solution to all nodes.

The aim of this work is to investigate tradeoffs between the round complexity for solving a task, and the bandwidth of the links. So far, results are known only for $B = O(\log n)$ in the classical CONGEST model, and $B = +\infty$ in the so-called LOCAL model. A specific case that deserves particular interest is to determine, given a task, the minimum value for B such that the task is solvable in $O(D)$ rounds in diameter-D networks with links of bandwidth B. For instance, in the case of the MST construction task, what is the minimum value of B

* Additional support from ANR Project DESCARTES.

licensed under Creative Commons License CC-BY
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

such that MST can be constructed in $O(D)$ rounds? It is known that $B = O(\log n)$ is not sufficient, as $\Omega(D + \sqrt{n})$ rounds is a lower bounds on the number of rounds for such a value of B [18].

In this paper, our objective is twofold. First, we are aiming at establishing tradeoffs between the size of the messages B and the number of rounds, by obtaining algorithms for the CONGEST_B model, having round complexities that are parametric on B. Our second objective is to better understand the role of B in the CONGEST_B model, and try to figure out if B always has some impact on the round complexity.

1.2 Our Results

Let $T_P(\mathcal{X})$ be the exact round complexity of a problem P using messages of size $B = O(\mathcal{X} \log n)$, given by ignoring additive factors that do not depend on B (e.g., the diameter). We define the speedup as $S_P(\mathcal{X}) = \frac{T_P(1)}{T_P(\mathcal{X})}$. Depending on S_P, we define three classes of problems and, for each of these classes, we provide examples of problems belonging to them.

- *bandwidth efficient*: the class of all problems P such that $S_P(\mathcal{X}) = \Theta(\mathcal{X})$, i.e., problems having a round complexity that fully scales with the bandwidth.
- *bandwidth sensitive*: this class contains all problems P such that $S_P(\mathcal{X}) = o(\mathcal{X})$ and $S_P(\mathcal{X}) = \Omega(1)$, i.e., those problems that have a round complexity that scales with the bandwidth, but not linearly.
- *bandwidth insensitive*: the class of all problems P such that $S_P(\mathcal{X}) = \Theta(1)$, i.e., whose complexity does not depend on the size of the messages.

First, we investigate the round complexity, in the CONGEST_B model, of the All Pairs Shortest Path (APSP) problem in unweighted graphs. We know that for this problem there is a lower bound of $\Omega(D + \frac{n}{B})$ rounds [6] and an algorithm performing in $O(n)$ rounds [11]. We show that the APSP problem is bandwidth efficient, i.e., that $S_{APSP}(\mathcal{X}) = \Theta(\mathcal{X})$, by modifying the existing algorithm to run in time $O(D + \frac{n \log n}{B}) = O(D + \frac{n}{\mathcal{X}})$.

We then investigate the round complexity of two well studied problems, namely Minimum Spanning Tree (MST) and Single Source Shortest Path (SSSP). For both these problems there is a lower bound of $\Omega(D + \sqrt{\frac{n}{B}})$ rounds [22]. On the other hand, these two problems have been studied only when $B = O(\log n)$. In this case, there exists an algorithm that solves the MST construction problem in $O(D + \sqrt{n} \log^* n)$ rounds [12], that is a round complexity that nearly matches the lower bound. For the SSSP problem, no sublinear algorithm that matches the lower bound and finds an exact solution is known. Conversely, we can find a $(1 + \epsilon)$-approximation of the solution in $\widetilde{O}(\epsilon^{-O(1)}(\sqrt{n} + D))$ rounds, that is optimal. We show that the round complexity of both these problems scales with B, that is, if it is possible to send messages of size B at each round,

- there exists an algorithm that constructs a MST in $\widetilde{O}(D + \sqrt{\frac{n}{B}})$ rounds;
- there esists an algorithm that finds a $(1+\epsilon)$-approximation of SSSP in $\widetilde{O}(\epsilon^{-O(1)}(\sqrt{\frac{n}{B}} + D))$ rounds.

Notice that these two round complexities match their respective lower bounds, for any value of B. Also, for both these problems, $S(\mathcal{X}) = \Theta(\sqrt{\mathcal{X}})$, i.e., these two problems are *bandwidth sensitive*.

We then show that there are problems for which, by increasing the bandwidth of the links, the round complexity does not change. In order to reduce the round complexity, one needs to increase the size of the messages exponentially. More specifically, we show that there is a problem, $Distance_k$, solvable in $O(\log n)$ rounds using messages of size $B = O(\log n)$, but such that if one wants to solve that problem in $\log n - 1$ rounds, then one needs to

use messages of size at least $\Omega(\frac{n}{\log^3 n})$. In other words, $Distance_k$ is a *bandwidth insensitive* problem for bandwidths in the range from $\Omega(\log n)$ to $O(\frac{n}{\log^3 n})$ bits, since in this range $S(\mathcal{X}) = \Theta(1)$.

2 Related Work

In general, the MST construction problem has been widely studied. In the distributed asynchronous context, Gallager, Humblet and Spira [7] gave an algorithm with time complexity $O(n \log n)$ that uses $O(|E| + n \log n)$ messages, which is optimal. In the synchronous setting, the first sublinear algorithm was given by Garay et al. in [8]. Its running time is $O(D + n^{\frac{\ln 3}{\ln 6}} \log^* n) \approx O(D + n^{0.61} \log^* n)$, where D is the diameter of the graph. This complexity was later improved to $O(D + \sqrt{n} \log^* n)$ in [12]. Then, Peleg et al. [21] showed that this latter complexity is nearly optimal, giving an $\Omega(D + \frac{\sqrt{n}}{\log n})$ lower bound, which was improved by Sarma et al. [22] to $\Omega(D + \sqrt{\frac{n}{\log n}})$ and then by Ookawa et al. [18] to $\Omega(D + \sqrt{n})$. All these lower bounds hold for graphs with diameter $\Omega(\log n)$. For constant diameter graphs, there are bounds of $\widetilde{\Omega}(n^{1/3})$ rounds for diameter 4, $\widetilde{\Omega}(n^{1/4})$ rounds for diameter 3, and $O(\log n)$ rounds for diameter 2 (see [15]). Then, Elkin [3] showed that, if termination detection is not required, the diameter of the graph is not a lower bound, and that there exists an algorithm that requires $\widetilde{O}(\mu + \sqrt{n})$ rounds, where μ is the so-called MST-radius of the graph. Recently, Pandurangan et al. [19] showed a randomized algorithm that is able to construct a MST in $\widetilde{O}(D + \sqrt{n})$ rounds while using an optimal number of messages, $\widetilde{O}(|E|)$. Elkin [5] showed how to achieve the same result deterministically. Notice that all the aforementioned algorithms use messages of size $B = O(\log n)$, while some of the lower bounds have been explicitly stated as a function of B. Among the various lower bounds, the best one that depends on B is the one of Sarma et al. [22], that is $\Omega(D + \sqrt{\frac{n}{B}})$. In the congested clique the MST problem can be solved much faster: Ghaffari et al. [9] gave an algorithm that solves the problem in $O(\log^* n)$ rounds using messages of size $O(\log n)$. They also showed that, by using slightly bigger messages, the complexity becomes $O(1)$ rounds.

Concerning the Single Source Shortest Path problem, an exact solution can be found using Bellman & Ford Algorithm in linear time. The problem of finding an exact solution in sublinear time remained opened for years, and only very recently Elkin [4] gave an algorithm that is able to solve the problem in $O((n \log n)^{\frac{5}{6}})$ rounds for $D = O(\sqrt{n \log n})$ and $O(D^{\frac{1}{3}} \cdot (n \log n)^{\frac{2}{3}})$ rounds for larger diameters. The problem of finding an approximate solution seems easier. In fact, Lenzen et al. [13] showed how to find an $O(\frac{1}{\epsilon})$-approximation in $\widetilde{O}(n^{\frac{1}{2}+\epsilon} + D)$ rounds. Then, Nanongkai [16] improved this result providing a randomized algorithm that finds a $(1+\epsilon)$-approximation in $\widetilde{O}(n^{\frac{1}{2}} \cdot D^{\frac{1}{4}} + D)$ rounds. Finally, Henzinger et al. [10] removed the multiplicative dependence on D, by giving a deterministic algorithm that solves the problem in $O(n^{\frac{1}{2}+o(1)} + D^{1+o(1)})$ and finds a $(1 + o(1))-$approximation. Becker et al. [2] improved this complexity by showing how to find a $(1+\epsilon)$ approximation in $\widetilde{O}(\epsilon^{-O(1)} \cdot (n^{\frac{1}{2}} + D))$ rounds. This last algorithm matches the lower bound given by Sarma et al. [22], that is $\Omega(D + \sqrt{\frac{n}{B}})$, but all the aforementioned algorithms are for $B = O(\log n)$.

The All Pairs Shortest Path problem is also well studied. Frischknecht et al. [6] showed that, in dense graphs, the diameter can not be computed in sublinear time (using small messages), by providing a lower bound of $\Omega(\frac{n}{B})$ rounds. This result implies a lower bound for the APSP problem as well. Then, Abboud et al. [1] provided an $\widetilde{\Omega}(n)$ lower bound for $B = O(\log n)$, even for sparse networks. Concerning upper bounds, Holzer et al. [11] showed how to solve the APSP problem deterministically, in $O(n)$ rounds, in unweighted graphs.

Nanongkai [16] presented a randomized algorithm that finds a $(1 + o(1))$-approximation in $\widetilde{O}(n)$ rounds in the weighted case. Lenzen et al. [14] showed that a $(1 + \epsilon)$-approximate solution can be deterministically found in $O(\epsilon^{-2} \cdot n \log n)$ rounds.

3 Model and Definitions

We are considering the classical CONGEST model for distributed network computing [20]. The network is modeled as a connected simple graph (no self-loops, and no parallel edges). The nodes of the graph are computing entities exchanging messages along the edges of the graph. Nodes are given arbitrary distinct identities (IDs) in a range polynomial in n, in n-node networks. Hence, every ID can be stored on $O(\log n)$ bits.

In the CONGEST model, all nodes start simultaneously, and execute the same algorithm in parallel. Computation proceeds synchronously, in a sequence of *rounds*. At each round, every node:

- performs some individual computation,
- sends messages to neighbors in the network, and
- receives messages sent by neighbors.

The main constraint imposed by the CONGEST model is a restriction of the amount of data that can be transferred between neighboring nodes during a round. The model in which the messages are limited to B bits is called CONGEST_B. Typically B is chosen to be $O(\log n)$, and CONGEST is typically used to refer to the $\text{CONGEST}_{O(\log n)}$ model.

The CONGEST model is well suited for analyzing the impact of limiting the throughput of a network on its capacity to solve tasks efficiently. The complexity of a distributed algorithm in the CONGEST model is expressed in number of rounds.

A similar model is the Congested Clique, where the communication graph is a clique and the same restrictions of the CONGEST model hold, i.e., all nodes can communicate with all the other nodes in one round using messages of size B. In the Broadcast Congested Clique each node is restricted to send the *same* message to all the other nodes.

4 All Pairs Shortest Path

In this section we show how to modify the APSP algorithm of [11] and reduce the round complexity when the bandwidth increases. Recall that, in the APSP problem, each node needs to find its distance from all the other nodes. In [11] is shown how to deterministically find an exact solution of the APSP problem in $O(D + n)$ rounds in unweighted graphs. The procedure is the following:

1. construct a BFS tree;
2. perform a DFS visit on the tree;
3. at each step of the DFS visit, wait 1 round and then start a BFS from the current node (if it is visited for the first time).

All nodes will know their distances from all the other nodes by knowing their distance from the root of each BFS tree. In [11] is shown that, by waiting one round before starting each visit, the breadth-first searches can be executed concurrently without congestion. In other words, at each round, at most one message passes on a fixed edge. The complexity comes by the fact that the DFS visit requires $O(n)$ rounds and that $O(D)$ is the time required to complete a BFS.

We show how to modify this procedure in order to solve the problem in $O(D + \frac{n \log n}{B})$ rounds.

▶ **Theorem 1.** *There exists an algorithm for the* CONGEST$_B$ *model that solves the All Pairs Shortest Path problem in* $O(D + \frac{n \log n}{B})$ *rounds.*

Proof. Let s_i, $1 \leq i \leq n$, be the i-th node visited (for the first time) by a DFS performed on the BFS tree. Let t_i, $0 \leq t_i < 2n$, be the time at which node s_i is visited *for the first time* by the DFS. Consider the sequence $\mathcal{T} = (t_1, \ldots, t_n)$. We split $\mathcal{T}$ in $\lceil \frac{B}{\log n} \rceil$ subsequences $T_1, \ldots, T_{\lceil \frac{B}{\log n} \rceil}$ of length at most $\lceil \frac{2\,n \log n}{B} \rceil$ each, where T_j contains the values t_i such that $\lceil \frac{2\,n \log n}{B} \rceil (j-1) \leq t_i < \lceil \frac{2\,n \log n}{B} \rceil j$. Notice that, if each node s_i starts its BFS at time $2\,t_i$, there is no congestion between different breadth-first searches, for the same arguments showed in [11]. Now, in order to speed up the computation, we can visit all the sequences concurrently, by starting the BFS of node s_i with $t_i \in T_j$ at time $2(t_i - \lceil \frac{2\,n \log n}{B} \rceil (j-1))$. In other words, we are splitting the DFS visit in at most $\lceil \frac{B}{\log n} \rceil$ parts, each of them of length at most $\lceil \frac{2\,n \log n}{B} \rceil$. By doing this, there can be congestion for at most $\lceil \frac{B}{\log n} \rceil$ breadth-first searches at the same time (one for each sequence), and since the bandwidth is B, a constant number of rounds is enough to transmit on the same edge $\lceil \frac{B}{\log n} \rceil$ messages of size $O(\log n)$ belonging to different breadth-first searches. Since each sequence is of length at most $\lceil \frac{2\,n \log n}{B} \rceil$, the total time required to solve the problem is $O(D + \frac{n \log n}{B})$.

What remains to show is how each node s_i can compute t_i in $O(D)$ rounds. Notice that, by performing a convergecast, all nodes of the BFS tree can learn the size of the subtree rooted at each of its children. Now, starting from the root s_1 and setting $t_1 = 0$, each node s_i can assign an arbitrary order to its children, and send, to each child s_j, t_i and the sum σ of the number of nodes present in the subtrees rooted on the children before s_j. Then each child s_j can compute $t_j = 2\sigma + t_i + 1$ and repeat the same procedure. In total, $O(D)$ rounds are required.

◀

5 Minimum Spanning Tree

In this section we show how to modify the MST construction algorithm of [12] and reduce the round complexity when the bandwidth increases, by proving the following theorem.

▶ **Theorem 2.** *There exists an algorithm for the* CONGEST$_B$ *model that constructs a Minimum Spanning Tree in* $\widetilde{O}(D + \sqrt{\frac{n}{B}})$ *rounds.*

Recall that, for the MST construction task, every node is given as input the weight $w(e)$ of each of its incident edges e. These weights are supposed to be of values polynomial in the size n of the network $G = (V, E, w)$, and thus each weight can be stored on $O(\log n)$ bits. The output of every node is a set of incident edges, such that the collection of all outputs forms an MST of the network. A Minimum Spanning Tree is a subset of edges $T \subseteq E$ such that (V, T) is connected and $\sum_{e \in T} w(e)$ is minimum. At the end of the computation each node must know which of its adjacent edges belong to the Minimum Spanning Tree.

The algorithm of [12] is divided in two phases: Fast-Dom-G and Pipeline. The first part, Fast-Dom-G, computes a set of sub-MST by applying a modified version of the algorithm of Gallager, Humblet and Spira (GHS) [7]. The second part, Pipeline, finds the remaining edges by performing a convergecast, where nodes send edges in non-decreasing order of weight by ignoring edges that close some cycle. Fast-Dom-G runs in $O(k \log^* n)$ rounds and creates $\frac{n}{k+1}$ fragments, where k is a parameter to be fixed later, while Pipeline runs in $O(D + \frac{n}{k})$ rounds. By setting $k = \sqrt{n}$, they obtain a complexity of $O(D + \sqrt{n} \log^* n)$ rounds.

We show that Pipeline can be performed faster if more bandwidth is allowed. More specifically, we show that it can be performed in $O(D + \frac{n \log n}{k \cdot B})$ rounds. By choosing $k = \sqrt{\frac{n \log n}{B}}$, we obtain Theorem 2.

▶ **Lemma 3.** Pipeline *can be completed in* $O(D + \frac{n \log n}{k \cdot B})$ *rounds, where* B *is the size of the messages.*

The rest of the section is used to prove Lemma 3, that is, we now show how to adapt the Pipeline procedure in such a way that if each node sends $O(\frac{B}{\log n})$ edges per round, its running is still fully pipelined. This implies that it is possible to send $O(\frac{n}{k+1})$ edges in $O(D + \frac{n}{k+1} / \frac{B}{\log n}) = O(D + \frac{n \log n}{(k+1)B}) = O(D + k)$ rounds.

The existing Pipeline algorithm, modified to send $\frac{B}{\log n}$ edges per round, does the following:

1. construct a BFS-tree of G;
2. each node knows, from the first phase (Fast-Dom-G), the edges that connect it to other fragments;
3. each node keeps a set of edges (in non-decreasing order of weight) not already sent to its parent;
4. leaves start at round 0, intermediate nodes start when they have received at least a message from all their neighbors;
5. at each step each node sends the $\frac{B}{\log n}$ lightest edges that does not close a cycle with the edges already sent (or the ones that it is sending) and repeats this step until there are no valid edges (since $\frac{B}{\log n}$ could be non-integral, we can use two rounds to complete this task);
6. at the end the root chooses the $\frac{n}{k+1}$ lightest received edges and broadcasts them.

We adapt the proof in [12] to show that the Pipeline procedure described above correctly completes within the required time. We first show the main part of the proof of [12] and then our adaptation.

In [12] is shown inductively that, at each step, there is at least one edge that can be upcasted. Let a node be *active* if it has sent an edge in the previous step. Assume that there is a node v that is active from m steps and it has at least one active child u. Notice that its active children are active from at least $m + 1$ steps. Let U be the set of edges already sent from node v to its parent in the previous steps. We know that this set forms a forest of trees, that we call $U_1, ..., U_l$. Let $x_i = |U_i|$. Notice that $\sum x_i = m$, since node v sent an edge in all the previous m rounds. We also know that $|V(U_i)| = x_i + 1$, since U_i is a tree. Let C be the set of edges that v already received from its child u. We know that $|C| \geq m + 1$, since each child is active from at least one step before its parent.

Now, suppose for the sake of contradiction that all the edges in C are not candidates that can be sent. It means that each edge, either it has already been sent, or that it closes a cycle. In both cases, the endpoints of each edge are part of some edge of U (the edges already sent). Since U forms a forest this also implies that there are no edges that connect two different trees U_i, U_j of U. It means that C can be partitioned depending on the trees to which the endpoints of each edge belong to. We can say that for each i, each edge in C_i has its endpoints in $V(C_i)$. And since C is a forest, we obtain that $|C_i| \leq |V(U_i) - 1|$. It follows that $|C_i| \leq |V(U_i) - 1| = x_i$, $C \leq \sum x_i = m \leq |C| - 1$ that is a contradiction. This implies that C contains at least one candidate.

The above proof can be adapted for our purposes in the following way. Since at each step a child sends $\frac{B}{\log n}$ edges, we know that $\sum x_i = |U| = \frac{mB}{\log n}$. We know also that children sent

edges for at least $m+1$ steps, thus $|C| \geq \frac{(m+1)B}{\log n}$. Again, for the sake of contradiction, assume that the number of candidate edges is less than $\frac{B}{\log n}$. This means that at least $|C| - \frac{B}{\log n} + 1$ edges have both their endpoints in U. In this case there are at least $|C| - \frac{B}{\log n} + 1$ edges that can be partitioned in sets $C_1, \ldots, C_l$ such that the endpoints of edges in C_i are only in $V(U_i)$. Hence, we obtain that $\sum |C_i| \leq \sum(|V(U_i)| - 1) = \sum x_i$. Since $\sum |C_i|$ is at least $|C| - \frac{B}{\log n} + 1$, we have that $|C| - \frac{B}{\log n} + 1 \leq \sum x_i$, implying that $|C| \leq \sum x_i + \frac{B}{\log n} - 1$. Thus, we obtain $\frac{(m+1)B}{\log n} \leq |C| \leq \sum x_i + \frac{B}{\log n} - 1 = \frac{mB}{\log n} + \frac{B}{\log n} - 1 = \frac{(m+1)B}{\log n} - 1$, which is a contradiction. This proves that C contains at least $\frac{B}{\log n}$ candidates.

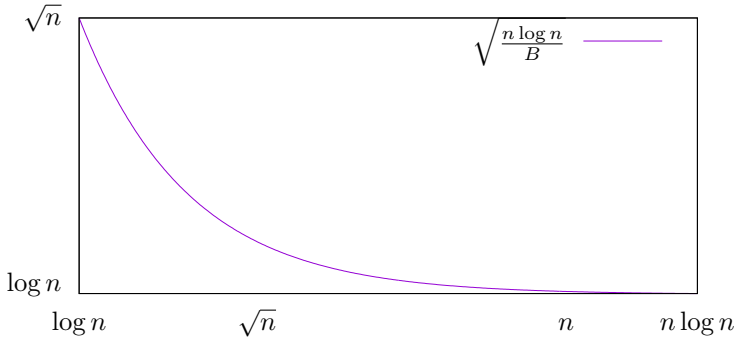

Figure 1 Round complexity as a function of B

6 Single Source Shortest Path

In this section we show how to modify the SSSP algorithm of [2] and reduce the round complexity when the bandwidth increases. Recall that in the SSSP problem, given a node v, all nodes of the graph have to find their distance from v. In [2] it is shown how to deterministically find a $(1+\epsilon)$ approximate solution of the SSSP problem using $\widetilde{O}(\epsilon^{-O(1)}(\sqrt{n}+D))$ rounds, where D is the (hop) diameter of the network. First, it exploits a Theorem of [10], that states the following:

▶ **Theorem 4.** *([10]) Given any weighted undirected network $G = (V, E, w)$ and a source node $s \in V$, there is a $\widetilde{O}(\sqrt{n})$-round deterministic distributed algorithm in the broadcast* CONGEST *model that computes an overlay network $G' = (V', E', w')$ with edge weights $w' : E' \to \{1, \ldots, poly(n)\}$ and some additional information for every node with the following properties:*

- *$|V'| = \widetilde{O}(\epsilon^{-1}\sqrt{n})$ and $s \in V'$.*
- *For $\epsilon' := \epsilon/7$, each node $v \in V$ can infer a $(1+\epsilon)$-approximation of its distance to s from $(1+\epsilon')$-approximations of the distances between s and each $t \in V'$.*

By applying Theorem 4, the SSSP problem is reduced to a graph of roughly $\sqrt{n}$ nodes (*skeleton* nodes). Then, in [10] authors provide an algorithm that solves the SSSP problem in the Broadcast Congested Clique and that runs in $\epsilon^{-9}polylog(n)$ rounds. They also show how to emulate this algorithm in the CONGEST model, with a slowdown proportional to D and to the number of nodes of the clique. Since the problem is reduced to an instance with $\sqrt{n}$ nodes, the algorithm can be emulated in $\epsilon^{-9}polylog(n) \cdot O(D + \sqrt{n})$ rounds. At this point the skeleton nodes can broadcast their distance from s to all the nodes of the original

graph in $O(D+\sqrt{n})$ rounds, and all the nodes can compute their approximate distance from s (Theorem 4).

As stated in [10], Theorem 4 is a deterministic version of a more general Theorem stated in [16], where $|V'|$ can be of size $\widetilde{O}(\alpha)$ and the running time of their algorithm is $\widetilde{O}(\alpha + \frac{n}{\alpha} + D)$ and succeeds with high probability. The round complexity is proved by providing an algorithm that solves the bounded-hop multi-source shortest path problem in $\widetilde{O}(|V'|+h+D)$ rounds, where the sources are nodes in V' and the number of hops h is $\frac{n \log n}{\alpha}$.

We now show that the aforementioned algorithm can complete in $\widetilde{O}(\frac{\alpha}{B} + \frac{n}{\alpha} + D)$ rounds, and that a round of communication of the broadcast congested clique can be emulated in $\widetilde{O}(\frac{|V'|}{B})$ in the CONGEST_B model, by proving the following lemmas.

▶ **Lemma 5.** *A round of communication of the Broadcast Congested Clique can be emulated in the* CONGEST_B *model in* $\widetilde{O}(\frac{|V'|}{B})$ *rounds.*

Proof. As in [2], we can solve the problem by using pipelining on a BFS tree. As in the MST case (Lemma 3), the speed of the pipelining linearly depends on $\frac{1}{B}$. ◀

▶ **Lemma 6.** *There is a distributed algorithm that runs in* $\widetilde{O}(\frac{|V'|}{B} + h + D)$ *rounds in the* CONGEST_B *model that solves the bounded-hop multi-source shortest path problem.*

Proof. The idea of the algorithm in [16] is to execute many bounded-hop single-source shortest path algorithms in parallel, one for each source node, and to randomly delay the starting time of each execution in order to avoid congestion. In [16] is shown that the execution of a single bounded-hop single-source shortest path algorithm requires $O(h + D)$ rounds using messages of logarithmic size, and by choosing a random delay from the interval of numbers from 0 to $|V'| \log n$, $|V'|$ executions can be performed in parallel without much congestion, obtaining a round complexity of $\widetilde{O}(|V'|+h+D)$. We show that, by increasing the bandwidth (thus allowing more congestion), it is possible to reduce the size of the interval from where the random delay is chosen.

Let $k = |V'|$. Let $\mathcal{M}_{i,u}$ be the set of messages sent by node u while executing the bounded-hop single-source shortest path algorithm for source s_i. In [16] (Lemma 3.7) it is shown that $\mathcal{M}_{i,u} \leq c \log n$ for some constant c. They then show that, if the delay is randomly chosen from the integers from 0 to $k \log n$, then the probability that there exists a time t, a node u and a set $\mathcal{M} \subseteq \bigcup_i \mathcal{M}_{i,u}$ such that $|\mathcal{M}| \geq \log n$ and all messages in $\mathcal{M}$ are broadcasted by u at time $t = O(k + h + D)$, is $O(\frac{1}{n^2})$. We extend this result showing in the remaining part of the section that, if the delay is randomly chosen from the integers from 0 to $\frac{k \log^2 n}{B}$, then the probability that $\mathcal{M} \geq B$ is also $O(\frac{1}{n^2})$, that implies the Lemma.

Fix any node u, time t and set $\mathcal{M}$ as above. Since u at each round sends at most one message for each s_i, we can assume that $|\mathcal{M} \cap \mathcal{M}_{i,u}| \leq 1$. This implies that $|\mathcal{M}| \leq k$. As in the original proof, we can bound the number of possible sets $\mathcal{M}$ of size m by $\binom{k}{m}(c \log n)^m$. This holds since each set $\mathcal{M}$ can be constructed by picking m different sets $\mathcal{M}_{i,u}$ and picking one message out of $c \log n$ messages from each $\mathcal{M}_{i,u}$. Then, the probability that all the messages of a set are sent at the same round is at most $(\frac{B}{k \log^2 n})^{|M|}$. Thus, for fixed u and t, the probability that there exists $\mathcal{M}$ such that $|\mathcal{M}| \geq B$ and all messages in $\mathcal{M}$ are sent by u at time t, is at most

$$\sum_{m=B}^{k} \binom{k}{m} (c \log n)^m (\frac{B}{k \log^2 n})^m \leq$$

$$\sum_{m=B}^{k} (\frac{ke}{m})^m (c \log n)^m (\frac{B}{k \log^2 n})^m =$$

$$\sum_{m=B}^{k} (\frac{ecB}{m \log n})^m \leq \sum_{m=B}^{k} (\frac{ecB}{B \log n})^B \leq$$

$$\sum_{m=B}^{k} (\frac{ec}{\log n})^B \leq k(\frac{ec}{\log n})^B$$

Since $B \geq \log n$, for large n the above formula is at most $\frac{1}{n^4}$. Then the lemma follows by summing this probability over all nodes u and time steps t. ◀

Now we have all the ingredients to prove the following theorem.

▶ **Theorem 7.** *There is a distributed algorithm that runs in $\widetilde{O}(\epsilon^{-O(1)}(\sqrt{\frac{n}{B}} + D))$ rounds in the CONGEST$_B$ model that solves the $(1+\epsilon)$ approximate SSSP problem.*

Proof. Let $|V'|$ be $\widetilde{O}(\epsilon^{-1}\sqrt{nB})$. By Theorem 4 and Lemma 6, we can reduce the problem to a SSSP instance on $|V'|$ nodes in $\widetilde{O}(\epsilon^{-1}\sqrt{\frac{n}{B}} + D)$ rounds. We can then simulate, using Lemma 5, the Broadcast Congested Clique algorithm of [10] in $\widetilde{O}(\epsilon^{-O(1)}\sqrt{\frac{n}{B}})$ rounds to solve the original problem. ◀

7 $Distance_k$

In this section we show a problem, called $Distance_k$, whose round complexity does not depend on the size B of the messages. Consider a graph $G = (V, E)$ that is the underlying communication graph, and consider a (possibly different) directed graph $G' = (V, E')$. To each node $v \in V$ is provided as input the set of its neighbors that it has in G' (its outgoing edges). Notice that the neighbors that v has in G and in G' could be different. The problem $Distance_k$ consists in finding a node $w \in V$ that is at distance k from a fixed node u in G'. A node w is at distance k from u if there is a directed path starting from u that ends at w and that path is the shortest one. An example of input instance of $Distance_k$ is shown in Figure 2, where, for example, a valid result for $Distance_2(1)$ is 7.

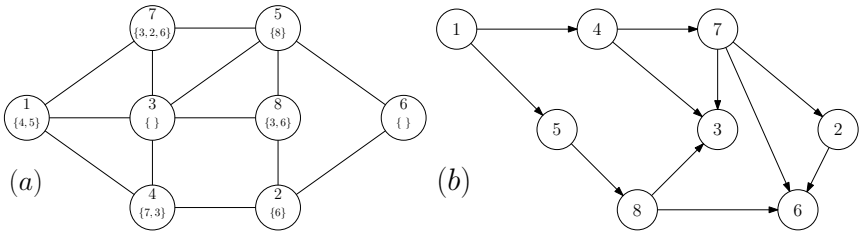

Figure 2 (a) represents G, the communication graph where the algorithm is executed and the input given to the nodes, while (b) is G', the directed graph given in input to the nodes.

▶ **Theorem 8.** *There are instances of $Distance_k$ that can be solved in $O(k)$ rounds using messages of size $B = O(\log n)$, that, if we want to solve in $k-1$ rounds, we need to use messages of size $B = \Omega(\frac{n}{k^3})$.*

In the following we will prove Theorem 8. Consider instances of $Distance_k$ where the out-degree of G' is at most 1. In this case, in the CONGEST model, using messages of size $O(\log n)$ bits, this problem can be solved in $O(k \cdot D(G))$ rounds, where $D(G)$ is the diameter of the communication graph, using Algorithm 1.

Algorithm 1 $Distance_k(u)$

```
1: if my Id is u then
2:     my distance is 0
3: else
4:     my distance is +∞
5: end if
6: for i ← 0, k − 1 do
7:     if my distance is i then
8:         broadcast in G the Id of my neighbor in G′
9:     end if
10:    if I receive a broadcast containing my Id ∧ my distance is +∞ then
11:        my distance is i + 1
12:    end if
13: end for
14: if my distance is k then
15:    broadcast my Id
16: end if
```

The idea is the following: once a node knows that it is at distance i, it can perform a BFS on G in order to broadcast its outgoing edge, that it has in G', to all the other nodes. Since each broadcast requires $O(D)$ rounds, in total $O(k \cdot D)$ rounds are required. Notice that, if the diameter of G is constant, the complexity becomes $O(k)$.

We now show that, in order to solve the problem in $k-1$ rounds, we need messages of size at least $\Omega(\frac{n}{k^3})$. Notice that, by choosing $k = \log n$, we obtain the following: there are instances of $Distance_k$ that can be solved in $O(\log n)$ rounds using messages of size $O(\log n)$, while in order to solve the problem in less than $\log n$ rounds, messages of size $\Omega(\frac{n}{\log^3 n})$ are needed, i.e., there is an exponential gap in the size of the messages.

In [17] the following problem, $Pointer_k$, is defined. There are two players, Alice and Bob, and to each of them is provided a list of n pointers, each pointing to the list of the other player. Their task is to follow these pointers, starting from some fixed pointer of Alice, and then find the k-th pointer. It is known that the communication complexity of this problem is $O(k \log n)$ bits: it is enough to send the i-th pointer at the i-th round, for a total of k rounds, where at each round $\log n$ bits are sent. They show that, in order to solve this problem in $k-1$ rounds, even if Las Vegas algorithms are allowed, the communication complexity of this problem becomes $\Omega(\frac{n}{k^2})$ bits. Notice that for $k = \log n$ there is an exponential gap between the two complexities.

We can reduce an instance of the problem $Pointer_k$ on p pointers to an instance of $Distance_k$ with $p+2$ nodes in the following way. Construct a graph $G = (V, E)$ where the nodes V are partitioned in two groups L and R of equal size, plus two special nodes l and r (see Figure 3). The nodes of each group are connected to a single special node and there is an edge connecting the two special nodes l and r. To each node of L is assigned a different ID from the set $\{1, \ldots, |L|\}$ and to each node of R is assigned a different ID from the set $\{|L|+1, \ldots, |L|+|R|\}$. We assign to node l the ID $|V|-1$ and to node r the ID $|V|$. The edges of $G' = (V, E')$ are defined in the following way. If the i-th pointer of Alice points to the j-th pointer of Bob, the node with ID i has an outgoing edge to the node with ID $|L|+j$. On the other hand, if the i-th pointer of Bob points to the j-th pointer of Alice, the node with ID $|L|+i$ has an outgoing edge to the node with ID j. Also, nodes

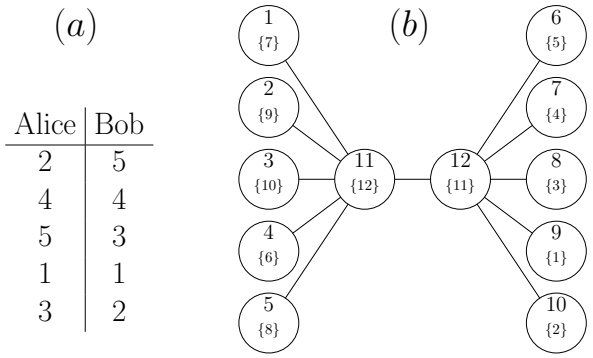

Figure 3 (a) shows the $Pointer_k$ instance, while (b) shows its $Distance_k$ version.

$|V| - 1$ and $|V|$ have an edge between themselves. An example of reduction is shown in Figure 3. Notice that Alice (resp. Bob) can construct the input of the nodes of L (resp. R) without communicating with Bob (resp. Alice). Assume that there is an algorithm for the CONGEST$_B$ model that solves the problem in $k-1$ rounds. Alice and Bob can simulate this algorithm by exchanging the B bits that are transmitted between the two special nodes at each round. Thus, a lower bound for $Pointer_k$ holds for $Distance_k$ as well.

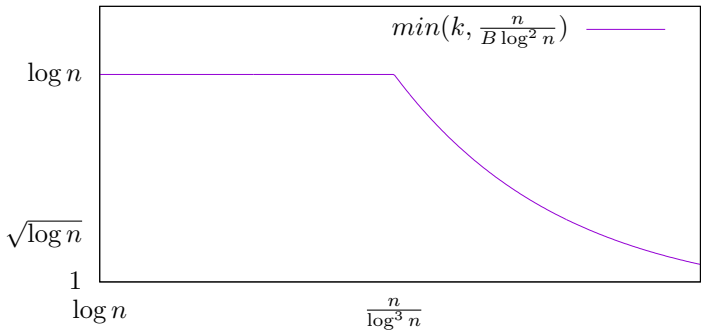

Figure 4 Round complexity as a function of B

In the CONGEST$_B$ model, we are allowed to send B bits on each edge at each round. If we want to solve $Distance_k$ in less than k rounds, we need to transfer $\Omega(\frac{n}{k^2})$ bits, thus $\Omega(\frac{n}{Bk^2})$ rounds are required. In other words, we need to satisfy $\frac{n}{Bk^2} < k$, that gives $B = \Omega(\frac{n}{k^3})$.

8 Conclusion

We considered the problem of understanding how the bandwidth affects the round complexity in the CONGEST model. Our results show that in some cases the round complexity perfectly scales with the bandwidth, matching existing lower bounds. On the other hand we showed that there are problems in which we need to increase exponentially the size of the messages in order to speed up the round complexity. These results suggest that it is important to analyze algorithms in the CONGEST model by considering a wider spectrum of bandwidths. It would be interesting to study this tradeoff between bandwidth and round complexity for other problems present in the literature.

Acknowledgements: I would like to thank my supervisor, Pierre Fraigniaud, for fruitful ideas and comments about the topic of this paper.

References

1 Amir Abboud, Keren Censor-Hillel, and Seri Khoury. Near-linear lower bounds for distributed distance computations, even in sparse networks. In *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, pages 29–42, 2016.

2 Ruben Becker, Andreas Karrenbauer, Sebastian Krinninger, and Christoph Lenzen. Approximate undirected transshipment and shortest paths via gradient descent. *CoRR*, abs/1607.05127, 2016.

3 Michael Elkin. A faster distributed protocol for constructing a minimum spanning tree. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 359–368, 2004.

4 Michael Elkin. Distributed exact shortest paths in sublinear time. *CoRR*, abs/1703.01939, 2017.

5 Michael Elkin. A simple deterministic distributed MST algorithm, with near-optimal time and message complexities. *CoRR*, abs/1703.02411, 2017.

6 Silvio Frischknecht, Stephan Holzer, and Roger Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1150–1162, 2012.

7 Robert G. Gallager, Pierre A. Humblet, and Philip M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 5(1):66–77, 1983.

8 Juan A. Garay, Shay Kutten, and David Peleg. A sublinear time distributed algorithm for minimum-weight spanning trees. *SIAM J. Comput.*, 27(1):302–316, 1998.

9 Mohsen Ghaffari and Merav Parter. MST in log-star rounds of congested clique. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing, PODC 2016, Chicago, IL, USA, July 25-28, 2016*, pages 19–28, 2016.

10 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A deterministic almost-tight distributed algorithm for approximating single-source shortest paths. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 489–498, 2016.

11 Stephan Holzer and Roger Wattenhofer. Optimal distributed all pairs shortest paths and applications. In *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 355–364, 2012.

12 Shay Kutten and David Peleg. Fast distributed construction of small k-dominating sets and applications. *J. Algorithms*, 28(1):40–66, 1998.

13 Christoph Lenzen and Boaz Patt-Shamir. Fast routing table construction using small messages: extended abstract. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 381–390, 2013.

14 Christoph Lenzen and Boaz Patt-Shamir. Fast partial distance estimation and applications. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 153–162, 2015.

15 Zvi Lotker, Boaz Patt-Shamir, and David Peleg. Distributed MST for constant diameter graphs. In *20th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 63–71, 2001.

16 Danupon Nanongkai. Distributed approximation algorithms for weighted shortest paths. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 565–573, 2014.

17 Noam Nisan and Avi Wigderson. Rounds in communication complexity revisited. *SIAM J. Comput.*, 22(1):211–219, 1993.

18 Hiroaki Ookawa and Taisuke Izumi. Filling logarithmic gaps in distributed complexity for global problems. In *SOFSEM 2015: Theory and Practice of Computer Science - 41st International Conference on Current Trends in Theory and Practice of Computer Science, Pec pod Sněžkou, Czech Republic, January 24-29, 2015. Proceedings*, pages 377–388, 2015.

19 Gopal Pandurangan, Peter Robinson, and Michele Scquizzato. A time- and message-optimal distributed algorithm for minimum spanning trees. *CoRR*, abs/1607.06883, 2016.

20 David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, 2000.

21 David Peleg and Vitaly Rubinovich. A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. *SIAM J. Comput.*, 30(5):1427–1442, 2000.

22 Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. In *43rd ACM Symposium on Theory of Computing (STOC)*, pages 363–372, 2011.